

FNAL dCache Admin Notes I

Version 2.0

Created: 11 March 2005

Modified: 10 January 2007

Rob Kennedy

- 1) Terms and Actors
- 2) Diagram of a User Interaction
- 3) FNAL dCache Configurations – (*subject to change!*)
 - a,b) FNDCA, FNDCAT overall, door ports, sub-caches
 - c,d) CDFDCA, CDFDCAT .. overall information only
 - e) CMSDCA no longer covered here
- 4) Responsibilities – need to be reviewed though.

Terms and Actors: Dictionary

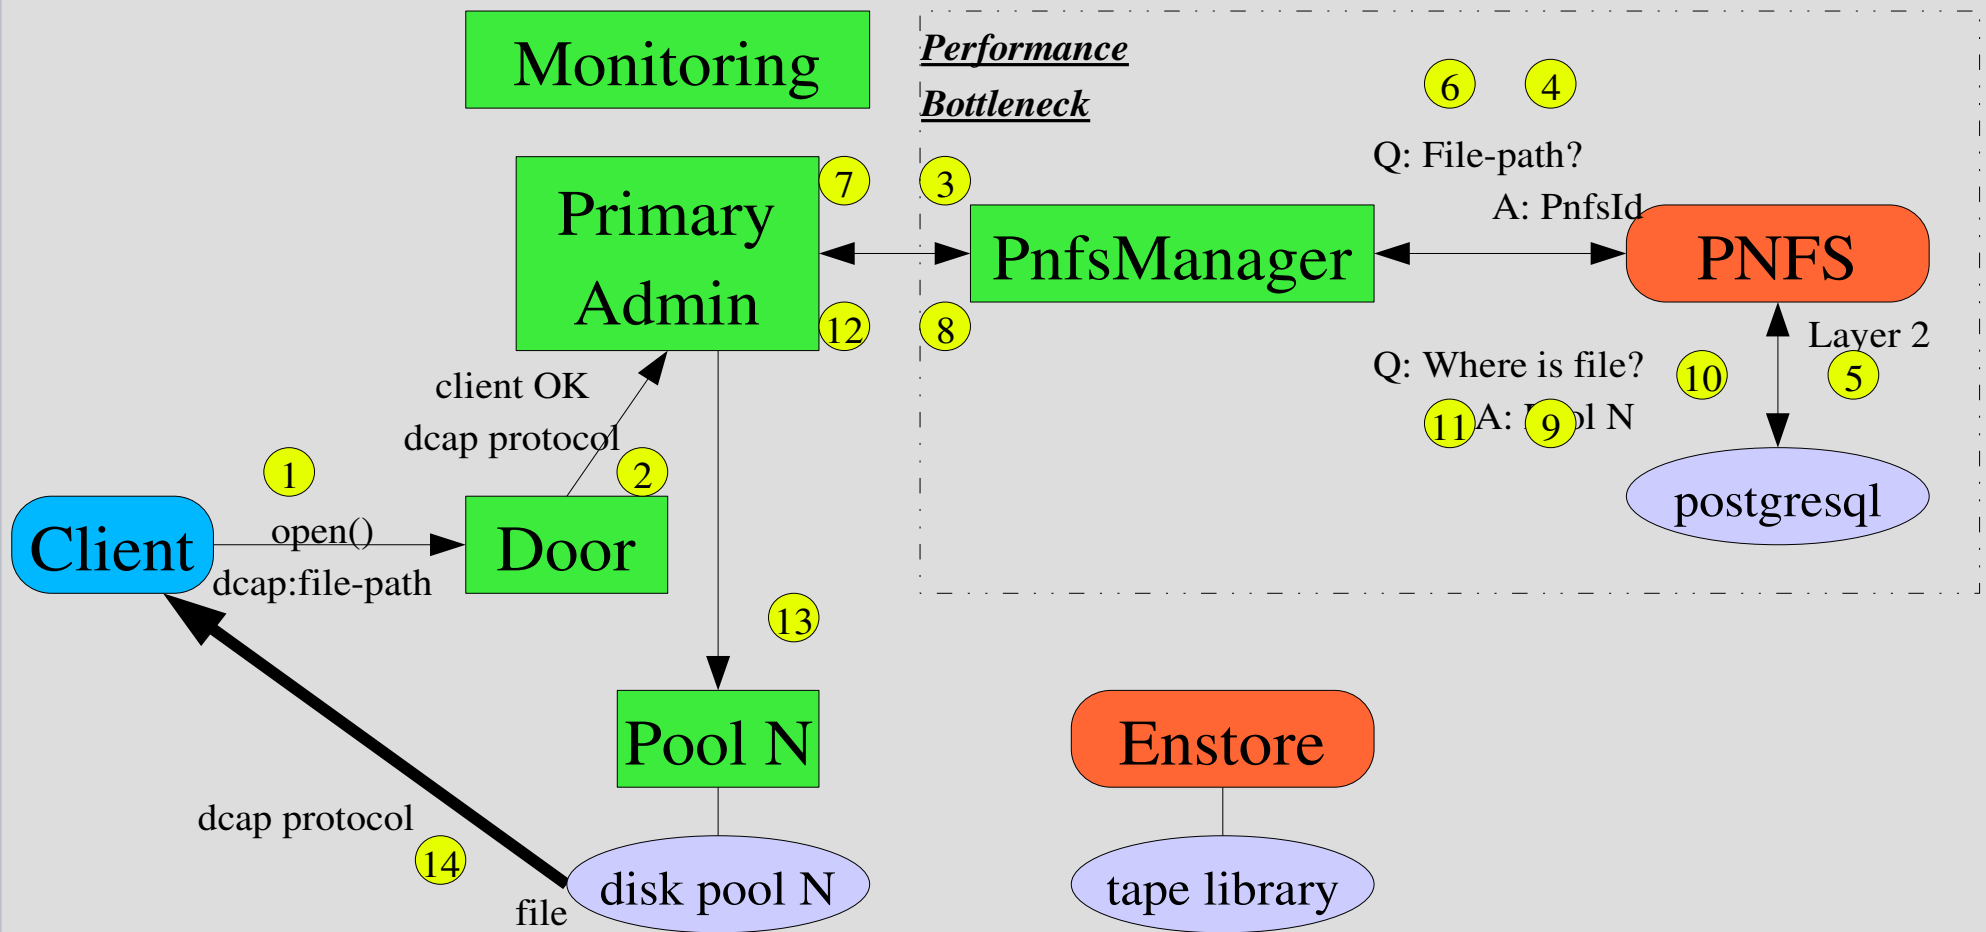
- Admin node: a host to dCache services
- Primary Admin node: SPOF host (single-instance services)
- Door node: a host to dCache “portals” or *doors* for clients
- Door: authenticate and authorize a client for a protocol
- Monitor node: a host performing various monitoring tasks
- Pool node: a host to dCache data file storage
- Pool: a virtual data partition in the dCache storage space
- Mover: process receiving or sending data to client.
- Restore: restoring data from MSS into cache.
- Store: storing data in cache to the MSS.
- P2P: pool-to-pool file transfers inside of a dCache system.
- PNFS: name space and replica catalog for dCache & Enstore

Terms and Actors: Service Structure

- Cell: smallest unit of dCache service. Each cell has the same infrastructure (cell nucleus), but supports varying functionality.
- Domain: a collection of cells running in one Java virtual machine. Cells in a domain not necessarily related in function.
- LocationManager: Singleton service. Knows about all domains, helps setup message *tunnels* amongst them. Abbreviated “lm”.
- PnfsManager: Singleton dCache *interface* to PNFS, is not the manager of PNFS. *Horribly confusing name!* dCache-PNFS interaction is known to be a performance bottleneck overall.
- SRM: Storage Resource Manager – Uniform Grid interface
- FNAL dCache: core dCache service, SRM, PNFS, FNAL-added monitoring plots, Dcap client library, and other misc.

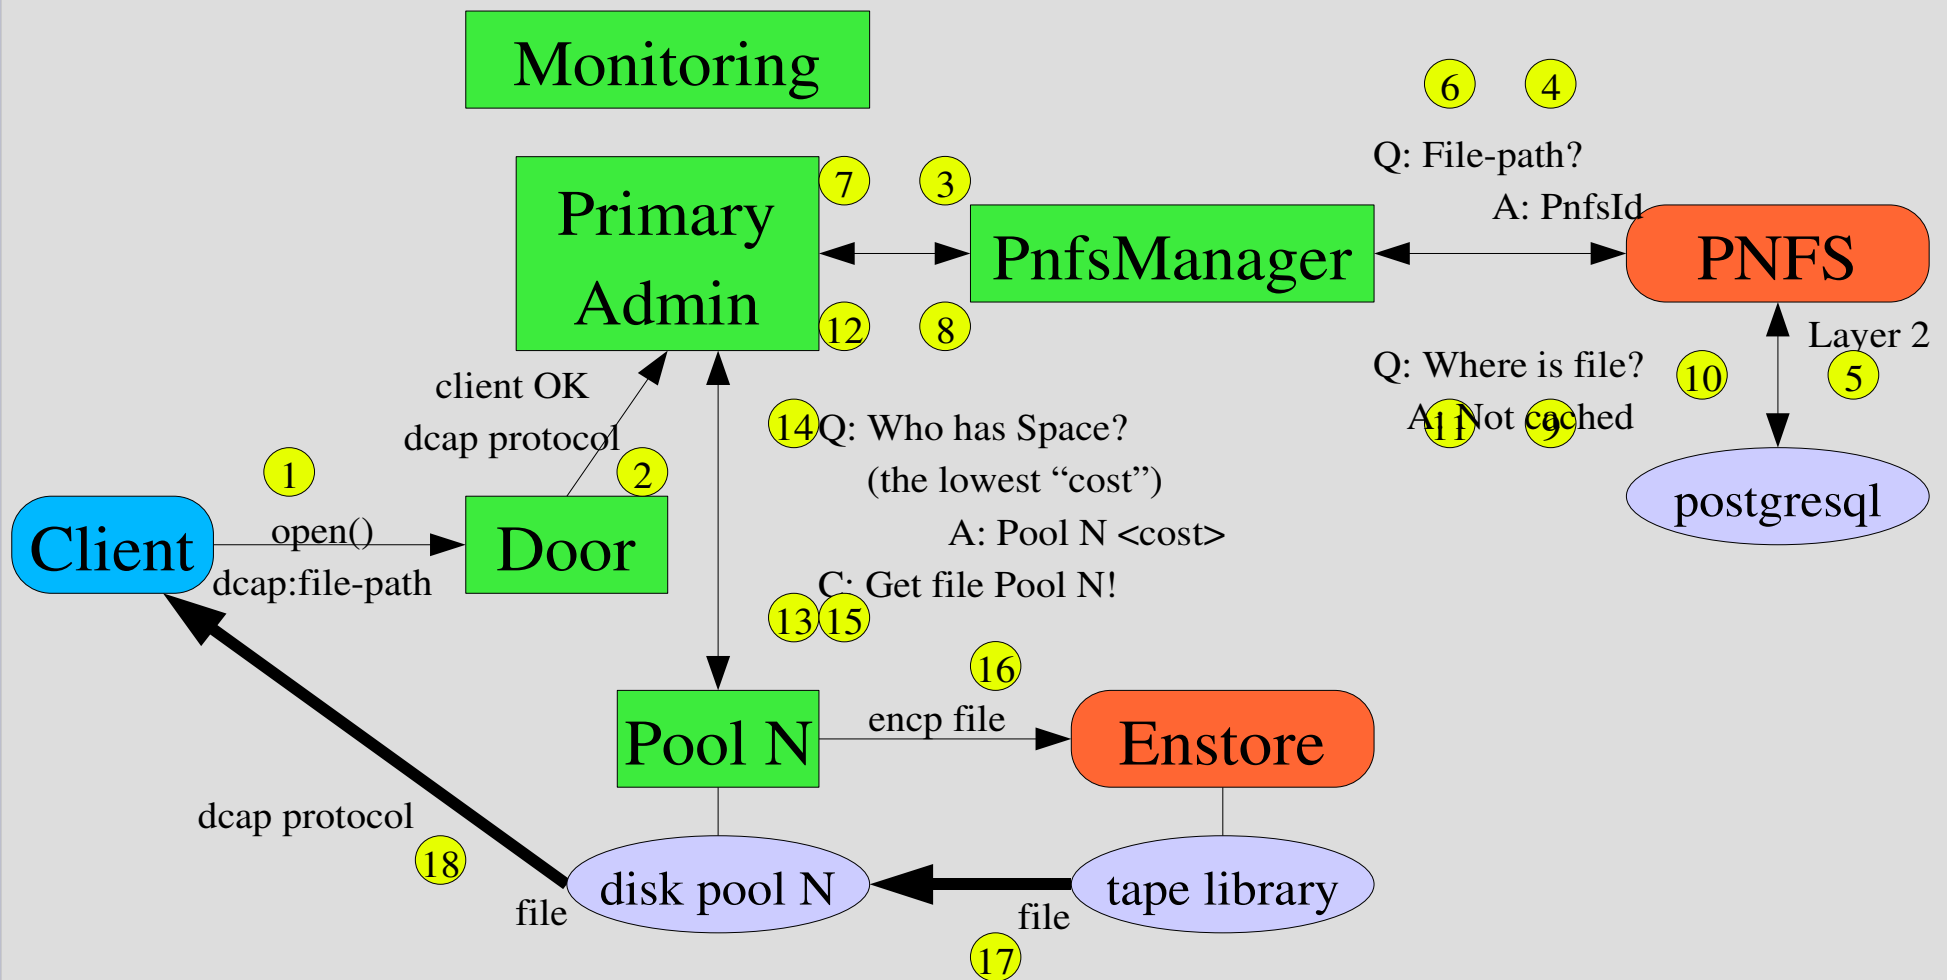
Client Reads a File in Cache

(simplified)



Client Reads a File not in Cache

(simplified)



FNDCA Configuration

Public dCache

- fndca (fndca3a): Primary admin node, web server host
- stkenstrv1: PnfsManager (“pnfsDomain”) Interface to PNFS.
- fndca1 (stkendca2a): 1 door node
- fndcam (fndca2a): monitor node
- Pool nodes (stkendca7a-20a): the file servers, 2-6 pools each.
- Doors: On door node at present (cell name):
 - 2 unsecured dcap (DCap00, DCap01) - read-only
 - 2 kerberized dcap (DCapK00, DCapK01)
 - 2 X509 dcap (DCapG00, DCapG01)
 - 1 weakly authenticated FTP (WFTP0) - read-only
 - 1 kerberized FTP (KFTP0)
 - 2 Grid FTP (GFTP0, GFTP1)
 - 1 SRM interface (SRM-stkendca2a)

FNDCA Doors

Public dCache

- | <u>DOOR TYPE</u> | <u>Start/stop Names</u> | <u>Host:port</u> |
|--|-------------------------|----------------------|
| • 2 unsecured dcap | door00, door01 | fndca1: 24125, 24136 |
| • 2 kerberized dcap | doorK00, doorK01 | fndca1: 24725, 24736 |
| • 2 X509 dcap | doorG00, doorG01 | fndca1: 24525, 24536 |
| • 1 weakly auth. FTP | weakftpd00 | fndca1: 24126 |
| • 1 kerberized FTP | kerberizedftpd00 | fndca1: 24127 |
| • 2 Grid FTP | gridftpd00, gridftpd01 | fndca1: 2811, 2812 |
| – <i>Note: 2811 is the standard gridftp port number. 2812 is a kluge</i> | | |
| • 1 SRM interface | <new deployment> | fndca1: 8443 |

FNDCA Sub-Caches

Public dCache

- “Sub-caches” are portions of a dCache system that exclusively serve a predefined set of file families for read or for write.
- “Link” names: <http://fndca3a.fnal.gov:2288/poolInfo/links/>*
- RawDataWrite: Used by Minos DAQ input, soon by CDMS too
 - 24 hour trigger before write to tape. Lots of reads permitted.
- KTeVRead: Used by KTeV
 - Generic read pool setup, restricted to KTeV file families.
- KTeVWrite: Virtual sub-cache for KTeV writes
 - Steers KTeV writes into the general write cache pools
- FermiGridVolatile: Used by FermiGrid users
 - Cache with no tape back-up. Short-term store for Farms jobs.
- read, write: General Read and Write caches... everything else.

FNDCAT Configuration

FNDCA Test dCache

- fndcat (stkendca3a): Primary admin, door, and monitoring node.
- PnfsManager on the head node. ONLY TRUE on test stands!
- Uses the same PNFS service as production FNDCA dCache.
- Pool nodes: stkendca6a + any pool nodes under test.
- Doors: the same as the production system, except:
 - All doors are on head node
 - There is not second gridftp door.
- Only general read and write caches are defined.
- The configuration of the test stand may vary to support tests.

CDFDCA Configuration

CDF Production dCache

- cdfdca (fcdfdcach1): Primary admin node
- cdfensrv1: PnfsManager (“pnfsDomain”) Interface to PNFS.
- cdfdca1,2,3 (fcdfdcach5,2,3): 3 door nodes
- cdfdcam (fcdfdcach4): monitor node
- Pool nodes: fcdfdca<nnn>
- Doors: On --each-- of the 3 door nodes:
 - 10 kerberized dcap (DCapK00-09,10-19,20-29)
 - 400 logins max per dcap door since on dedicated hosts.
 - 2 X509 dcap doors (DCapG00-01, 10-11, 20-21)
 - 1 WFTP<n>, 1 KFTP<n>, 1 GFTP<n>, where n=1,2,3
 - SRM interface active only on cdfdca1 (srm1).
- Several Sub-caches: http://cdfdca.fnal.gov:2288/poolInfo/links/*

CDFDCA Doors

CDF Production dCache

- | <u>DOOR TYPE</u> | <u>Start/stop Names</u> | <u>Host:port</u> |
|-------------------|------------------------------|---------------------|
| • Kerberized dcap | | |
| • doorK00-09 | cdafdca1: 25125, 25136-25144 | |
| • doorK10-19 | cdafdca2: 25145-25154 | |
| • doorK20-29 | cdafdca3: 25155-25164 | |
| • X509 dcap | | |
| • doorG00-01 | cdafdca1: 25525, 25536 | |
| • doorG10-11 | cdafdca2: 25545-25546 | |
| • doorG20-21 | cdafdca3: 25555-25556 | |
| • Kerberized FTP | kerberizedftpdoor0,1,2 | cdafdca1,2,3: 25127 |
| • Grid FTP | gridftpdoor0,1,2 | cdafdca1,2,3: 2811 |
| • SRM interface | srm1 | cdafdca1: 8443 |

CDFDCAT Configuration

CDF Test dCache

- cdfdcat (fcdfdcach6): Primary admin, door, and monitoring node.
- PnfsManager on the head node. ONLY TRUE on test stands!
- Uses the same PNFS service as production CDFDCA dCache.
- Pool nodes: fcdfdata112 + any pool nodes under test.
- Doors: the same as the production system, except:
 - All doors are on the one teststand head node
 - Door port numbers are prefixed by 26xxx instead of 25xxx
- Only general read and write caches are defined.
- The configuration of the test stand may vary to support tests.

Responsibilities

- FNDCA(T): SSA responsibility. Dcache developers on-call for consult.
 - 7x24x365 – head node, door node, overall service or data at risk
 - 7x8 best effort – individual node problem where data no at risk
-
- CDFDCA(T): sensitive and should be reviewed. Last understanding:
 - CDF-DH/Offline Ops: filters out user “non-dCache” requests.
 - H/w and OS – CDF buys. Run2-sys installs, prepares, and admins.
 - dCache operations – CDF-DH does, SSA helps, dCache devs consult
 - Example: CRC check failures – SSA treats.
 - dCache config – CDF-DH does most, dCache devs consult
-
- CMSDCA: SSA – just Cmspnfs2. Dcache devs often consult though.